National Science Foundation
Workshop

SIGNAL PROCESSING

Keystone, Colorado

November 15-17, 1991

Sponsored by

Circuits and Signal Processing Program

Microelectronics and Information Processing Systems Division

TABLE OF CONTENTS

Workshop Coordinator and Report Editor

Professor Delores M. Etter
Department of Electrical and Computer Engineering
University of Colorado, Boulder

1 EXECUTIVE SUMMARY

An NSF-supported workshop was held on November 15-17, 1991, at the Keystone Conference Center in Keystone, Colorado. One goal of the workshop was to evaluate the success of the last decade of research in signal processing by identifying important technological developments that depended on signal processing research results. Another goal was to provide a focus for the identification of important research topics with broad ranges of applications that are critical to the growth and competitiveness of technology. This focus also provides both a short-term vision and a long-term vision of the important areas and applications for signal processing research.

1.1 Participants

In order to accomplish the workshop goals, it was important that the participants represent a wide base of well-respected expertise from academia, industry, and research laboratories. A group of 28 participants were selected by the Workshop Coordinator; Appendix A contains a list of the participants and their affiliations. These participants were separated into four groups which represent the following research areas:

One-Dimensional Signal Processing
Multidimensional Signal Processing
Digital Representation
Implementation

Each group spent two days preparing a group report which is included as an appendix. In addition, a group of people who were unable to attend the workshop read the draft report and provided input for the final report. These readers are also listed in Appendix A.

1.2 General Observations

Novelty is an important part of university research, but the product is the key component in industry. It is important to learn to recognize and promote research results that can provide a real advantage to an application product. Encouraging this transition from research results to industrial implementation is critical.

The committee also identified several areas of research which they believed would have a significant and immediate positive effect on theoretical and practical signal processing work. These include the establishment of some sort of community-wide, easily accessible library or repository for contributed signal data, and a library of software and algorithms “known to work,” such as the old IEEE/ASSP DSP tapes. These, it was believed, would improve the effectiveness of many research and development efforts by removing the need to code and test “building block” algorithms used in the study of more advanced schemes.

Signal processors have often done badly in the public relations game and have lost out to other fields in the press (and hence, as a result, sometimes with Congress). The study group agreed that communication or “outreach” to others will significantly contribute to the continued growth of the signal processing field. This outreach must be both “outward” in the sense of communicating useful signal processing ideas and techniques to other areas, such as the manufacturing industry, and “inward” in the sense of seeking out expertise in cognate disciplines such as mathematics and physics.

1.3 Findings and Recommendations

The findings and recommendations addressed the following topics:

1. Significant technological advances that include signal processing components.

2. Important areas of signal processing research and their applications.

3. Areas for special focus.

Each group made a number of observations and several recommendations. A group of four findings and recommendations are presented that emerged from more than one group and generated strong support during the final workshop session.

1.3.1 Finding and Recommendation #1

Finding #1:

The signal processing community has not done a good job of informing the key leaders in government, industry, and academia of the many ways that signal processing is used in the technological advances of the last decade. This misperception of the significance of signal processing is detrimental to signal processing research in times of limited research funding.

Recommendation #1:

This workshop report should have wide distribution in order to help inform key leaders in government, industry, and academia about the signal processing components of technology for today and for tomorrow. The report should also be distributed to new faculty who may need guidance in selecting research topics and applications in areas of national significance.

1.3.2 Finding and Recommendation #2

Finding #2:

Manufacturing is an especially relevant topic to technology. Signal processing can play an important part in this evolving area, from communications on the workfloor, to non-destructive testing, to the design of robotics with vision and speech capabilities, to environmental sensors. Implementation of signal processing techniques is as important as the theoretical issues in this application.

Recommendation #2:

Researchers in signal processing should join forces with industry to identify research that will benefit manufacturing. We encourage the National Science Foundation to hold a workshop to bring members of the manufacturing industry together with signal processing researchers to facilitate this interaction and to develop a synergy in manufacturing research.

1.3.3 Finding and Recommendation #3

Finding #3:

Signal processing results provided the basis for many major medical breakthroughs, especially in the area of medical imaging. Significant results in signal processing algorithms, digital representations, and implementations should have a substantial impact in the area of biotechnology, an area of increased national emphasis. The interpretation and understanding of the initial results of the human genome project is another new area to apply signal modeling and information extraction techniques. Signal processing is also a key element in developing better sensory and computational aids for the handicapped.

Recommendation #3:

Cooperative efforts between the medical/biomedical community and researchers in signal processing should be encouraged by NSF and others. This interaction should include not only imaging diagnosis and aids for the handicapped but also interactions with molecular biologists to investigate modeling biological phenomenon such as the genome.

1.3.4 Finding and Recommendation #4

Finding #4:

The area of multimedia communications (interactive communications between human and computer involving several media and multiple senses) is an area with immense economic opportunity and very challenging signal processing problems. In addition, significant signal processing results are needed in the area of personal communication networks. Both areas require research results that allow for more efficient representation and that make use of efficient and effective algorithms to reduce bandwidth and provide robust decoding. Related areas are image fusion and image understanding.

Recommendation #4:

We strongly recommend that the National Science Foundation develop a formal research initiative in the area of signal processing applications to multimedia communication and personal communication networks. Joint projects should be encouraged with other programs at the National Science Foundation such as the Networking and Communications Program and programs in the Electrical Engineering Division.

1.4 Conclusions

Signal processing is an essential component in providing the national competitiveness needed in important application areas such as manufacturing, biotechnology, and communications. It is critical that research in the general area of signal processing, and specifically in the areas mentioned in the findings and recommendations, be continued. The results of signal processing research have already made a far-reaching impact on technology, and this research will continue to improve our competitive edge.

2 INTRODUCTION

The following definition of signal processing was developed at a meeting in 1991 of the National Science Foundation's MIPS (Microelectronics and Information Processing Systems) Advisory Committee:

> *Definition*: Signal processing is the extraction of information-bearing attributes from measured data, and any subsequent transformation of those attributes for the purposes of detection, estimation, classification, or waveform synthesis.

At the same Advisory Committee meeting, the signal processing members recommended that a workshop be held to provide both a short-term vision and a long-term vision.

Listing the important topics in signal processing is difficult because the area of signal processing has many facets, such as analysis, modeling, synthesis, algorithm development, and tool development. Signal processing also includes a number of applications areas such as communications, speech processing, medical imaging, video processing, and all areas of sensing. Signal processing algorithms are often implemented using critical technologies, such as high performance computers, application-specific integrated circuits (ASICs), very large scale integration (VLSI), and multiple-chip modules (MCMs).

This workshop was designed with several specific objectives in mind. First, an enumeration of the successes of signal processing research is required to support the claim that signal processing has been an integral part of many very important technological developments of the last decade. Secondly, a focus for the identification of the important research topics with broad ranges of applications is needed. The participants were selected to provide a wide-base of expertise from academia, industry, and research laboratories. The selection of the participants was determined by the Workshop Coordinator.

A group of 28 participants were invited. The sessions at the workshop were divided into the four areas listed below; the group coordinator's names are included:

- One-Dimensional Signal Processing
  - John Treichler, Applied Signal Technology, Inc.
- Multidimensional Signal Processing
  - David Munson, University of Illinois
- Digital Representation
  - Robert Gray, Stanford University
- Implementation

Earl Swartzlander, University of Texas, Austin

The group coordinators acted as moderators for the group meetings, and were responsible for summarizing the group discussions in the final workshop session. The group coordinators were also responsible for preparing the group reports contained in the appendices of this report.

3 GROUP REPORT SUMMARIES

The following summaries give a brief overview of the group reports. More details are included in the complete reports in the appendices.

3.1 One-Dimensional Signal Processing

This group focused almost entirely on one-dimensional (1-D) digital signal processing (DSP), bypassing traditional analog signal processing and newer areas such as optical signal processing. While they felt that these omissions were justified by the flexibility, programmability, precision, and repeatability of DSP implementations, they also pointed out that continued research is important in analog areas including analog ASIC design and optical signal processing.

The "enabling technologies" that have contributed to the high degree of research success in signal processing were identified and appear in the complete report. The committee also identified several areas where they believed that a moderate investment of energy and money could have a significant immediate positive effect on theoretical and practical signal processing work. These include the establishment of some sort of community-wide, easily accessible library or repository for contributed signal data, and a library of software and algorithms "known to work," such as the old IEEE/ASSP DSP tapes. These, it was believed, would improve the effectiveness of many R&D efforts by removing the need to code and test "building block" algorithms used in the study of more advanced schemes.

The theoretical areas which were identified as being especially important in one-dimensional DSP include representations and models (multi-scale, eigenspace, non-linear dynamical), algorithms and architectures, optimization and approximations, and information extraction (such as time-frequency methods, subspace methods, higher-order statistics, inverse methods, multi-sensor and data fusion).

A number of current and future applications of 1-D DSP were identified. These belong to topic areas which include real-time signal processing and imaging, communications, scientific computation and instrumentation, manufacturing, and medical signal/image processing. The study group agreed that communication or "outreach" to others will significantly contribute to the continued growth of the signal processing field. This outreach must be both "outward" in the sense of communicating useful signal processing ideas and techniques to other areas, such as the manufacturing industry, and "inward" in the sense of seeking out expertise in cognate disciplines such as mathematics and physics.

Finally, this study group selected two "problems" that it agreed were worthy of high-profile initiatives. Each needs signal processing ideas and technology as a basis for acceptable solutions. The areas recommended for special focus are

- SIGNAL PROCESSING IN EFFICIENT RF SPECTRAL UTILIZATION
  - voice coding and processing
  - channel coding
  - other digital communications technologies
- SIGNAL PROCESSING IN MANUFACTURING AND PRODUCT QUALITY CONTROL
  - nondestructive testing
  - shop communications
  - long-term product monitoring (location, status)

3.2 Multidimensional Signal Processing

Multidimensional signal processing (MDSP) involves the acquisition, manipulation, and display of multidimensional data. MDSP has close ties to a number of other disciplines, including communications, control systems, computer vision, computer graphics and visualization, computation, and VLSI. Also, by its very nature, MDSP has drawn on basic theory from mathematics and physics, and combined this theory with signal processing techniques in a wide range of applications areas. These include a variety of applications which fall under the general categories of imaging, image communication, image understanding, applications to basic sciences, and array processing. As examples, medical imaging is an evolving field with recent significant accomplishments that include helical-scan and dynamic 3-D computer tomography, phased-array ultrasound, and magnetic resonance imaging (MRI). In the area of remote sensing, synthetic aperture radar (SAR) and radio astronomy have aided in the development of stereo, polarimetric, and interferrometric imaging, and the mapping of the surface of Venus.

Important research needs in MDSP that were delineated include topics oriented toward specific problem areas such as higher-dimensional image processing, content-addressable image databases, image quality evaluation and theoretical bounds, and image fusion, and topics of a more theoretical nature such as nonlinear processing, multiscale and time-frequency analysis, optimization techniques, and algebraic methods.

Finally, the committee felt that there are two interdisciplinary research areas involving considerable MDSP. It can be expected that many technological spin-offs will result from new research initiatives in these areas. The areas recommended for special focus are

- MULTIDIMENSIONAL DIGITAL SIGNAL PROCESSING IN MULTIMEDIA

    - techniques for handling large amounts of video data
    - multidimensional compression
    - image understanding

- MULTIDIMENSIONAL DIGITAL SIGNAL PROCESSING IN INTELLIGENT VEHICLE HIGHWAY SYSTEMS

    - aerial SAR for traffic monitoring
    - automotive safety imagery
    - traffic light control
    - road surface monitoring

3.3 Digital Representation

Digital representation is the result of converting an input signal into digital data, or the starting point for synthesizing an output signal. The goal of digital representation is efficient (low cost, low delay, low power, simple algorithms, suitability for subsequent processing) and effective representation (high quality, preserves important information). Digital representation can be subdivided into analog-to-digital conversion (ADC), compression, decomposition, and modeling. ADC is primarily an implementation issue, with tradeoffs between analog and digital signal processing. Compression relates to issues such as perceptual coding and robust techniques. Also included are multiresolution techniques such as subband coding, quadrature mirror filters (QMF), wavelets, and trees. Modeling includes feature selection and learning theory.

The group identified several specific research areas, both practical and theoretical, which they believe merit increased effort and support. Experimental areas included the development of better perceptual models, and model verification and extension for data analysis. Theoretical areas included efforts to determine the behavior of signal decompositions such as quadrature mirror filter banks and wavelet decompositions in the presence of quantization noise and aliasing, and the development of perceptually meaningful distortion measures.

The group also discussed several “random topics” of interest including the whole issue of standards, university/industrial relations, tools, and education. They also provided a number of very timely general comments which get to the heart of many of the problems which face the signal processing community. For example,

> Novelty is the key in university research, but the product is king in industry. It is important to learn to recognize research results and promote research efforts that can provide a real advantage to an application product. Perhaps a separate research funding apparatus should be developed for university research that focuses more on short- and long-term development of new techniques that can be transferred to industry for implementation.
>
> EE signal processors have often done badly in the public relations game and have lost out to other fields in the press (and hence, as a result, sometimes with Congress).
>
> American academic engineering has a reputation for publishing theories, while the Japanese turn those theories into products.

This group also delineated a number of "hot" applications which they felt were of unusual importance to U.S. industry as well as being likely to benefit from advances in signal processing:

- entertainment: digital audio, TV, and multimedia
- medical image processing
- medical data analysis
- information services, databases
- personal communications
- metrology
- nondestructive testing
- manufacturing processes/digital process control
- automotive electronics
- sensor (and other) data transmission, storage, reconstruction
- voice and image synthesis
- system identification (for adaptation, adaptive equalization)
- digital representations that support editing and composition
- remote education (HDTV)
- environmental monitoring
- voice mail

Finally, three application areas stood out as being particularly important for U.S. industry. All three are commercially important, can significantly benefit from advances in signal processing, and can provide excellent training in a broad range of signal processing and related disciplines. The areas recommended for special focus are

- MULTIMEDIA

    - combining voice video, images, data
    - ability to telecommute

- PERSONAL COMMUNICATIONS

    - communication to the person, not the location
    - coding, compression, networking

- ENTERTAINMENT

    - video and audio
    - advanced television

3.4 Implementation

Implementation integrates disparate work in algorithms, architectures, and technologies in the context of specific applications. Examples of DSP that have only recently become possible due to new implementation techniques include ultrasonic doppler for cardiac blood flow imaging, cellular phones, synthetic aperture radar, and automobile sensors. Important areas that need new implementation methods include speech recognition, active noise control, and HDTV.

The recommendations of this group centered on a two-pronged approach.

> First, develop a framework for signal processing system design. This structure provides an environment for a unified and interrelated set of signal processing tools ranging from high-level functional specification to detailed physical implementation of the technology.
>
> Periodically exercise the design tools with exciting new applications which have significant signal processing challenges.

Research which leads to or in some way supports the development of signal processing system design tools should be supported. These design tools should specifically seek to overcome the barriers which traditionally separate algorithm, architecture, and technological pursuits. Because of the immediate importance of signal processing to many different application areas and the universal need for rapid prototyping, there should be a specific emphasis on performance directed synthesis. In addition, the research should be performed in the context of a specific application.

The group also cited a variety of research opportunities:

- new concurrent topologies and algorithms
- new, general algorithmic representational formalisms
- mapping techniques between different algorithmic structures
- automatic scheduling and resource allocation techniques (compilers) for parallel and/or deeply pipelined systems
- effective library generation and utilization techniques for algorithms, schedules, cells, software, etc.
- signal processing-oriented debugging and simulation tools.

Finally, the implementation group recognized that the implementation of a complete design system will require a considerable effort from the community at large. Such a system would consist of a common design infrastructure on which individual research tools can be developed and distributed. The consensus was that this infrastructure will be as important to the signal processing community as MOSIS has been to the VLSI community. The area recommended for special focus is

- SIGNAL PROCESSING INFRASTRUCTURE
  - graphical and textual entry
  - abstraction of dataflow, structure, physical technology
  - tools to transform between levels and within levels

4 CONCLUSIONS AND RECOMMENDATIONS

Four specific recommendations were selected as representative of research common to several areas. These areas also received strong support in the final workshop session.

4.1 Finding and Recommendation #1

Finding #1:

The signal processing community has not done a good job of informing the key leaders in government, industry, and academia of the many ways that signal processing is used in the technological advances of the last decade. This misperception of the significance of signal processing is detrimental to signal processing research in times of limited research funding.

Recommendation #1:

This workshop report should have wide distribution in order to help inform key leaders in government, industry, and academia about the signal processing components of technology for today and for tomorrow. The report should also be distributed to new faculty who may need guidance in selecting research topics and applications in areas of national significance.

4.2 Finding and Recommendation #2

Finding #2:

Manufacturing is an especially relevant topic to technology. Signal processing can play an important part in this evolving area, from communications on the workfloor, to non-destructive testing, to the design of robotics with vision and speech capabilities, to environmental sensors. Implementation of signal processing techniques is as important as the theoretical issues in this application.

Recommendation #2:

Researchers in signal processing should join forces with industry to identify research that will benefit manufacturing. We encourage the National Science Foundation to hold a workshop to bring members of the manufacturing industry together with signal processing researchers to facilitate this interaction and to develop a synergy in manufacturing research.

4.3 Finding and Recommendation #3

Finding #3:

Signal processing results provided the basis for many major medical breakthroughs, especially in the area of medical imaging. Significant results in signal processing algorithms, digital representations, and implementations should have a substantial impact in the area of biotechnology, an area of increased national emphasis. The interpretation and understanding of the initial results of the human genome project is another new area to apply signal modeling and information extraction techniques. Signal processing is also a key element in developing better sensory and computational aids for the handicapped.

Recommendation #3:

Cooperative efforts between the medical/biomedical community and researchers in signal processing should be encouraged by NSF and others. This interaction should include not only imaging diagnosis and aids for the handicapped but also interactions with molecular biologists to investigate modeling biological phenomenon such as the genome.

4.4 Finding and Recommendation #4

Finding #4:

The area of multimedia communications (interactive communications between human and computer involving several media and multiple senses) is an area with immense economic opportunity and very challenging signal processing problems. In addition, significant signal processing results are needed in the area of personal communication networks. Both areas require research results that allow for more efficient representation and that make use of efficient and effective algorithms to reduce bandwidth and provide robust decoding. Related areas are image fusion and image understanding.

Recommendation #4:

We strongly recommend that the National Science Foundation develop a formal research initiative in the area of signal processing applications to multimedia communication and personal communication networks. Joint projects should be encouraged with other programs at the National Science Foundation such as the Networking and Communications Program and programs in the Electrical Engineering Division.

All areas of technology today use signal processing. The applications range from adaptive coding of speech signals, to analysis of medical images, to echo-canceling in cellular telephones. The successful implementation of signal processing algorithms in both software and hardware is a result of more than a decade of research into various aspects of signal processing, including analysis, modeling, synthesis, algorithm development, tool development, and implementation. It is vital to the technological development of our country that research in signal processing research have high priority in research funding. This report recommends that specific attention be given to areas of manufacturing, biotechnology, and communications. New initiatives in the areas of multimedia communications and personal communications should be developed.

APPENDIX A PARTICIPANT LIST

One-Dimensional Signals and Systems

C. S. Burrus	(Rice Univ.)
Delores M. Etter	(Univ. of Colorado)
Mostafa Kaveh	(Univ. of Minnesota)
C. L. Nikias	(USC)
A. J. Paulraj	(Stanford)
William T. Rhodes	(Univ. of Colorado)
Louis L. Scharf	(Univ. of Colorado)
John R. Treichler	(Applied Signal Technology, Inc.)

Multidimensional Signals and Systems

Jan P. Allebach	(Purdue)
Michelle Covell	(SRI)
T. S. Huang	(Univ. of Illinois)
Russell M. Mersereau	(Georgia Tech)
David C. Munson, Jr.	(Univ. of Illinois)
A. Murat Tekalp	(Univ. of Rochester)
Sally L. Wood	(Santa Clara Univ.)

Digital Representation

Philip Chou	(Xerox Parc)
Robert M. Gray	(Stanford)
N. S. Jayant	(AT&T)
Tom D. Lookabough	(Compression Labs, Inc.)
Stephen A. Townes	(MITRE)
Donald W. Tufts	(Univ. of Rhode Island)
Martin Vetterli	(Columbia)

Implementation

Jonathan Allen	(MIT)
Thomas P. Barnwell III	(Georgia Tech)
Robert W. Broderson	(Berkeley)
Leah H. Jamieson	(Purdue)
Teresa Meng	(Stanford)
Malcolm G. Slaney	(Apple Computer)
Earl E. Swartzlander, Jr.	(Univ. of Texas Austin)

Readers

Faye Boudreaux-Bartels	(Univ. of Rhode Island)
John M. Cioffi	(Stanford)
George Cybenko	(Univ. of Illinois)
Dan E. Dudgeon	(MIT/Lincoln Labs)
Ronald L. Fante	(MITRE)
J. L. Flanagan	(Rutgers)
Mary J. Irwin	(Penn State)
T. Kailath	(Stanford)
Jae S. Lim	(MIT)
Dick Lyon	(Apple)
S. L. Marple, Jr.	(ORINCON)
James H. McClellan	(Georgia Tech)
Alan V. Oppenheim	(MIT)
Sarah A. Rajala	(NC State)
Ronald W. Schafer	(Georgia Tech)
John W. Woods	(Rensselaer)

NSF

Aubrey Bush	(Networking and Communications)
John Cozzens	(Circuits and Signal Processing)
John Lehmann	(MIPS)

APPENDIX B ONE-DIMENSIONAL SIGNAL PROCESSING REPORT

This report summarizes the discussions, conclusions, and recommendations of the One-Dimensional Signal Processing Group of the NSF Signal Processing Workshop held at Keystone, Colorado, 15-17 November, 1991. The group members were C. S. Burrus, D. M. Etter, M. Kaveh, C. L. Nikias, A. J. Paulraj, W. T. Rhodes, L. L. Scharf, and J. R. Treichler.

B.1 Background and Overview

Signal processing has been instrumental in the development of a variety of technologically important and financially successful products ranging from medical imaging equipment to cellular telephones and compact disc players. These developments directly exploited the body of theoretical signal processing knowledge built up over the last several decades. It can be reasonably expected that the number of new applications of this technology will grow and that extensions of this theory will permit the development of solutions for even more practical problems.

As part of a more general examination of the field of signal processing, a study group was formed to focus on the subfield of one-dimensional signal processing. This document constitutes a summary of the observations made by the study group and certain recommendations made by it.

For reasons outlined in Section B.2, this report focuses virtually completely on digital signal processing, bypassing traditional analog signal processing and newer areas such as optical signal processing. Furthermore, the principal focus of the reported deliberations was on the theoretical underpinnings of signal processing and not on its application or mechanization. Section B.3 contains a dissection of the key aspects of signal processing theory and identifies certain of these as ones of particular interest. As indicated in Section B.3, some of these are so identified because they are new and appear to be rich in possible theoretical results, while others are so identified since more understanding of those topics would facilitate practical application. As will be repeated there, *the identification of these areas is not intended to indicate that the others listed nearby are not of value but rather that the ones indicated merit special attention*.

Signal processing is a theoretical discipline that solves problems found in industrial, defense, and scientific applications. Its high degree of success in these applications stems in significant measure from the concomitant evolution of microelectronics. This and other so-called “enabling technologies” are also expected to be crucial to the continued practical application of signal processing. A list of those thought to be the most important over the next decade is provided in Section B.4.

The noted success of signal processing in practical applications is documented to some degree in the list appearing in Section B.5, accompanied by some very advanced potential applications. While the focus of this report is on the theoretical and scientific underpinnings of signal processing, it is important to observe that applications and possible applications have always been the instigation for many areas of analytical inquiry and will continue to be. In this light, it is necessary for theoreticians to be conversant with these practical problems and to understand the implications of these applications to the evolving theoretical base.

The study group's consensus is that communication or “outreach” to others will significantly contribute to the continued growth of the signal processing field. As elaborated in Section B.6, this outreach should be both “outward” in the sense of communicating useful signal processing ideas and techniques to other areas, such as the manufacturing industry, and “inward” in the sense of seeking out expertise in mathematics and physics.

The final section of this report recommends two specific areas for high-profile, well-funded initiatives. In each case, the initiative depends on signal processing for its successful achievement, addresses an area important to the public interest, and its achievement will in turn enrich the signal processing field's body of theoretical knowledge.

B.2 Does Digital Signal Processing Have a Future and Why?

The work of the “1-D study group” focused virtually entirely on digital signal processing as opposed to the variety of analog signal processing methodologies available. In the collective view of the study group, this focus is justified by the following points.

- DSP implementations feature flexibility, programmability, precision, and repeatability in degrees much greater than in analog signal processing (ASP) implementations.

- The performance and cost effectiveness of DSP implementations have been and will continue to be enhanced by “enabling technologies” such as microelectronics--much more so than analog implementations.

- Mathematical expressions can be implemented directly and with little approximation with DSP while ASP does not allow this, leading to the ability to use far more sophisticated algorithms with DSP than ASP would allow.

- Digital hardware can usually be multiplexed in a far more straightforward way than can analog hardware, providing opportunities for significant improvements in cost effectiveness.

These points, coupled with the expectation that they will remain true over the next decade, argue for substantial continued investment in digital signal processing theory and in the technologies which will enable its implementation. In passing, however, it should be noted that analog signal processing is not moribund. It is still the preferred implementation technology for a range of applications for which DSP is not yet suitable, including those requiring very high bandwidths, very high dynamic ranges, very low power, very low cost, or a high degree of parallelization. Continued investment is therefore appropriate here as well, notably in the areas of analog ASIC design and optical signal processing.

B.3 Theoretical Foundations

The signal processing community has developed a rich body of theoretical knowledge over the past two decades. The simultaneous evolution of microelectronics has permitted a very high “conversion rate;” that is, many theoretical concepts have already been applied to hardware or software which is successful in the commercial, scientific, and defense marketplace. This high degree of success, and the general belief that future theoretical work will also be successfully applied, motivated the “1-D” study group to carefully examine the theoretical framework of one-dimensional digital signal processing with the goal of identifying areas which might be particularly fruitful. To accomplish this, a general taxonomy of the field was produced and then specific areas were highlighted.

A taxonomy of one-dimensional signal processing can be done in a variety of ways. The study group chose to break the field down in a way similar to a typical decomposition of a signal processing system. In system design it is common to first determine a representation or model for the signals to be processed. The signals are processed using algorithms implemented in some form of hardware or software architecture. The algorithms are usually designed to achieve some type of optimization, often using approximations of signals, systems, or objective functions to make the analytical problem tractable. The intent of the signal processing system is usually to extract information, the exact type depending upon the application at hand. Since this particular decomposition of signal processing systems is commonly used and understood in both theory and practice, it serves as the basis for the next several subsections.

B.3.1 Representations and Modeling Theory

- wavelets (e.g., subband coding) and multiscale representations
- Gabor and other cellular stationary spectral representations (e.g. Fourier analysis)
- nonstationary spectral representations (e.g., time/frequency distributions)
- Markovian state/space representations (e.g., ARMA modeling)
- eigenspace (subspace) and principal components (e.g., Karhunen-Loeve)
- modal representations
- low-rank linear algebraic models
- nonlinear dynamic models
- set membership models
- fuzzy sets
- high-order statistical representations
- chaotic models and fractals
- point processes
- infinite-variance processes
- generalization to space-time vector, and multidimensional problems
- quantized and discretized representations
- partial differential equation models (e.g., field equations)
- physically accurate models
- transducer and measurement-matched models
- spline models

Coalesced "focus" areas
- multiscale representations
  (including wavelets, Gabor, spline, subband, . . .)
- eigenspace, modal, and principal components
  (including low-rank, linear algebraic models)
- fractal and higher-order statistical representations
- nonlinear dynamic systems (chaos, fractals, fuzzy, set membership)

B.3.2 Algorithms and Architectures

The identification of certain areas with an asterisk indicates that they merit special attention.

Algorithms
- * transforms
- * nonlinear adaptive systems
- * numerical linear algebra (including tensors)

- fast algorithms
- filtering
- acquisition and tracking

Architectures

* orthogonal
* filter banks
* multirate
- single/multiple-input, multiple data (SIMD/MIMD)
- lattice
- pipelined 1-D and 2-D
- dataflow

* Mapping algorithms to architectures
    - index mapping and matrix factorization
    - VLSI issues
    - hardware description languages/mapping tools

B.3.3 Optimization and Approximation for Signal Processing

The identification of certain areas with an asterisk indicates that they merit special attention.

* nonlinear (NL) optimization
- IQ (iterative quadratic)
- extremal problems
* l_1, l_{INF}, and l_α
- reduced-rank approximations
- dynamic programs
- annealing (stochastic and otherwise)
- homotopy
- non-simplex linear programming
* minimization of arithmetic operations
- genetic algorithms
- quadratic programming
- numerical linear algebra
- projections and convex sets
* discretization and quantization
- matrix perturbation theory

B.3.4 Information Extraction

- detection and classification (multistatic detection)
- estimation
- time series analysis

- array processing (multisensor time series analysis)
- spectrum analysis
- frequency-wavenumber spectrum analysis
- algorithms based on time-frequency distributions
- Fourier transforms
- short-term Fourier, Wigner-Ville
- quantization
- low-rank modeling
- wavelet transforms
- subspace methods
- matched field processing
- linear prediction
- data adaptive linear prediction
- data adaptive matched subspace processing
- adaptive filters
- measurement design
- Bayesian inference
- dynamic programs
- higher-order statistics
- data fusion (signal and data fusion)
- modal analysis

Coalesced "high-impact" areas

- methods using higher-order statistics
- time/frequency methods
- information extraction from NL dynamic (chaotic) systems
- inverse problems
- multisensor methods (e.g., array processing)
- subspace methods
- signal and data fusion

B.3.5 Established Disciplines

One disadvantage of the "system design" decomposition listed above is that it tends to hide the fact that established disciplines exist within the signal processing community. These disciplines have often grown up focused around a particular class of applications and problems. While the application objectives are different, all of these disciplines are active and vital, each draws upon the elements above, and all have made contributions to the understanding of those elements. A noninclusive list of these established disciplines and some of the outstanding problem areas for each includes the following:

filter design
complex approximation for IIR filters

    - rapid design
    - nonlinear (median, stack, . . .)
    - design of multirate and time-varying filters
- array processing
    - beamforming
        - multipath-rich environments
        - wideband
        - computational improvements
    - matched-field processing
        - “robust beamforming”
    - frequency / wavenumber analysis
    - DOA estimation
        - experimental data sets
        - validation of estimation
        - calibration
- spectrum analysis
    - frequency / wavenumber maps
    - nonstationary
- speech processing
- signal compression
- inverse problems
- blind deconvolution / signal equalization

B.4 Enabling Technology

Signal processing is a theoretical discipline that solves problems found in industrial, defense, and scientific applications. Its high degree of success in those applications stems in significant measure from the concomitant evolution of digital microelectronics. This and the other “enabling technologies” listed in Table B.4.1 are expected to be crucial to the continued practical application of signal processing over the next decade. A clear implication is that these enabling technologies also need support and development energy as well in order to maximize the beneficial effects of efforts applied to the development of signal processing concepts.

Table B.4.1: Technologies which Enable Signal Processing Success

- computer-aided design (CAD) tools
* application software and libraries
* signals database
- VLSI (microelectronics)
- test and developmental tools (e.g., oscilloscopes for complex-valued digital signals)
- memory technology (for computing in general, and to permit memory/computation tradeoffs)
- materials and devices beyond silicon
- hardware and software testbeds
- software tools and workstations
- computer architecture

The items marked with an asterisk are those thought by the committee to be areas in which a moderate investment of energy and money could have a significant immediate positive effect on theoretical and practical signal processing work. Obtaining actual signal data to validate theoretical results is a recurring problem within the university and corporate research environment. The establishment of some sort of community-wide, easily accessible library or repository for contributed signal data could largely ameliorate this problem. Similarly, a library of software and algorithms “known to work,” such as the old IEEE/ASSP DSP tapes, would improve the effectiveness of many R&D efforts by removing the need to code and test “building block” algorithms needed for the study of more advanced schemes. It should be emphasized that such libraries not only improve the productivity of university, government, and corporate research efforts but also allow commercial operations to develop DSP-based products without heavy investment in research staff and facilities.

The committee's consensus was that Optical Signal Processing (OSP) will grow significant in theoretical and practical interest in the medium and long terms. In the near term, OSP can be viewed as an enabling technology to more conventional analog and digital signal processing. The principal research interest should be in the areas of finding ways to trade the high bandwidth and parallelism available with OSP to improve its numerical and implementation accuracy. Attainment of accuracies typical of DSP hardware will be necessary for its widespread application.

B.5 Applications - Current and Future

B.5.1 Real-Time Signal Processing and Imaging

scanning electron microscope (SEM) imaging
ultrasonic
morphological
scanning instruments
entertainment (multimedia)
HDTV
medical
CAT (Computerized Axial Tomography)
NMR (Nuclear Magnetic Resonance)

B.5.2 Communications

cellular mobile phone
voice and pen computing
FAX
business
personal communication network (PCN)
wireless InterNet (business, shop floor)
data storage (magnetic, optical)
very small aperture terminal (VSAT)
global positioning system (GPS) navigation
control
HDTV
speech recognition, speech coding, and image compression
equalization
encryption, secure communications

B.5.3 Scientific Computation and Instrumentation

digital oscilloscopes, test equipment
partial differential equation solutions
inverse problems
visualization

optimization (algorithms and architectures)
spectrum analysis
crystallography
astrophysical imaging
seismic
tomography
nondestructive testing (NDT)
experimental design

B.5.4 Manufacturing

mask design
chemical rates of reaction
chemical quality control
aluminum smelting
voltage-current images
manufacturing fault analysis
robotics
inspection (NDT)
instrumentation
diagnostics (signature analysis)
active noise control

B.5.5 Medical Signal/Image Processing

digital hearing aids
implantable devices (e.g., smart defibrillator, drug release)
imaging
EKG (electrocardiogram), EEG (electroencephalogram) monitoring
ultrasound
diagnostics
aids for vision
compression for medical signals/imaging
smart prostheses

B.6 Outreach

The study group agreed that communication or “outreach” to others will significantly contribute to the continued growth of the signal processing field. This outreach must be both “outward” in the sense of communicating useful signal processing ideas and techniques to other areas, such as the manufacturing industry, and “inward” in the sense of seeking out expertise in mathematics and physics.

Table B.6.1 Outward and Inward Outreach.

Outward:

- education - spreading the word about signal processing concepts and techniques to other scientific and engineering fields and to those in the commercial sector to permit product development
  - traditional university education
  - workshops and short courses
  - provide signal processing tools
- cooperation with manufacturing industry

Inward:

- collaboration with mathematicians and scientists
  - SP community needs more mathematical knowledge to address nonlinear SP problems, e.g. chaotic systems
- device physics
- pressing problems spawn new ideas, e.g. measurement and sensing problems in environmental science

B.7 Recommended Initiatives

The extraordinarily rapid growth of the signal processing field is a testament to the practical success of products and systems which use the technology. Because of this success, there is no doubt that talent and resources will continue to be applied to SP-based products in ever greater measure over the next decade. Even so, there are still application areas in which significant problems remain to be solved. Some of the problems are technical, some are political, and some are caused by the perception that the technical problems can't be solved. From a number of suggestions, this study group selected two “problems” that it agreed were worthy of high-profile, well funded initiatives. Each needs signal processing ideas and technology as a basis for acceptable solutions, but the ultimate solution will no doubt be driven by many other considerations as well. The solution of each is in the public interest, and the achievement of success in both areas will in turn enrich the signal processing field's body of theoretical knowledge.

- SIGNAL PROCESSING FOR EFFICIENT RF SPECTRAL UTILIZATION
  - in voice processing
  - in voice coding
  - in channel coding

- in digital communications

- SIGNAL PROCESSING FOR MANUFACTURING AND PRODUCT QUALITY CONTROL

    - nondestructive testing
    - shop communications
    - long-term product monitoring (location, status)

APPENDIX C MULTIDIMENSIONAL SIGNAL PROCESSING REPORT

This report summarizes the discussions, conclusions, and recommendations of the Multidimensional Signal Processing Group of the NSF Signal Processing Workshop held at Keystone, Colorado, 15-17 November, 1991. The group members were J. P. Allebach, M. Covell, T. S. Huang, R. M. Mersereau, D. C. Munson, A. M Tekalp, and S. L. Wood.

C.1 Introduction

Multidimensional signal processing (MDSP) comprises the acquisition, manipulation, and display of multidimensional data. In this section of the report we concentrate primarily on evolving topics in MDSP, with an emphasis on imaging systems and image processing. Some important areas within MDSP, such as image compression and array processing, are dealt with primarily elsewhere in this report. We do not discuss, at all, classical multidimensional filter design and implementation, and traditional transforms, since these are relatively mature topics of study.

MDSP has close ties to a number of other disciplines, including communications, control systems, computer vision, computer graphics and visualization, computation, and VLSI. Also, by its very nature, MDSP has drawn on basic theory from mathematics and physics, and combined this theory with signal processing techniques in a wide range of application areas.

C.2 Applications

Important applications of MDSP include those shown in Table C.1. Medical imaging is a huge and still expanding field, which, although originated primarily by physicists, is now progressing with the help of researchers in MDSP. Recent, significant accomplishments in this field include helical-scan and dynamic 3-D computer tomography, phased-array ultrasound, and dramatic improvements in magnetic resonance imaging (MRI). Each of these accomplishments involves sophisticated signal processing. Some of the modalities for medical imaging are also used in general nondestructive testing. In the area of remote sensing, synthetic aperture radar (SAR) and radio astronomy have been spectacular successes. Recent accomplishments in SAR include the development of stereo, polarimetric, and interferrometric imaging, and the mapping of the surface of Venus. MDSP is finding increasing use in other areas of imaging, such as in crystallography, electron microscopy, and confocal microscopy, for the inspection of materials and the analysis of biological specimens. In the consumer arena, low-resolution, all-electronic cameras already have been introduced, and we are on the verge of a complete transformation of the photographic industry as we know it. MDSP will play a prominent role in this revolution.

The field of image communication is growing rapidly, and, to date, we have seen only the tip of the iceberg. The CCITT (Comite Consultatif Internationale Telegraphique et Telephonique) standard for FAX permitted the rapid infusion of document transmission into our daily lives. New standards are now nearly finalized for the transmission of digital video. We can expect digital electronic communication to revolutionize the office workplace, the publishing industry, and the entertainment industry. We are already witnessing 64 Kb/s video conferencing, and the future points to multimedia with integrated speech, text, and video, and the prospects for virtual reality. Progress in these areas will rely heavily on MDSP.

TABLE C.1. APPLICATIONS OF MDSP

I. IMAGING

- A. Medical Imaging and Nondestructive Evaluation
  1. X-Ray
  2. X-Ray Computer Tomography (CT)
  3. Magnetic Resonance Imaging (MRI)
  4. Ultrasound
  5. Single Photon Emission Computer Tomography (SPECT)
  6. Positron Emission Tomography (PET)
  7. Impedance
- B. Remote Sensing
  1. Radar Imaging
  2. Acoustic Imaging
  3. Infrared, Ultraviolet
  4. Radioastronomy
- C. Microstructure Imaging
  1. X-Ray Crystallography
  2. Electron Microscopy
  3. Confocal Microscopy
- D. Consumer Imaging
  1. Electronic Photography
  2. Advanced Video Cameras

II. IMAGE COMMUNICATION

- A. Electronic Publishing & Document Processing
- B. Advanced Television

C. Videoconferencing
D. Multimedia
   1. 3-D Models, Graphics for Visualization
   2. Virtual Reality - training (e.g., flight simulation), entertainment, architecture, interactive video aspects

III. IMAGE UNDERSTANDING

A. Manufacturing & Robotics
B. Security
C. Intelligent Vehicle Highway Systems (IVHS)
D. Agriculture

IV. APPLICATIONS TO BASIC SCIENCES

Archeology, Biology, Fluid Dynamics, Kinesiology, Meteorology, etc.

V. ARRAY PROCESSING

A. Source Localization
B. Seismic Exploration
C. Active Noise Control

The topics of image analysis and understanding are being actively applied in the fields of manufacturing and robotics, and security. Although this research has traditionally fallen within the province of researchers in computer vision, we are coming to discover that many problems in these areas can profitably benefit from combined approaches using ideas from both computer vision and MDSP. Future application areas that are expected to benefit from a combined approach range from agriculture to intelligent vehicle and highway systems. In addition, there are emerging applications of MDSP and computer vision in basic sciences. For example, we may wish to visually identify, and subsequently control, eddies in turbulent fluid flow, or to analyze motion of the human skeletal system.

Array processing is an important technique that has long been used for accurate target and source localization. The subject of array processing is considered more fully elsewhere in this report, but we do want to indicate here that new array processing techniques are being applied very successfully to the problems of seismic exploration and active noise control.

C.3 Future Research

Important research needs in MDSP include topics oriented toward specific problem areas, and also topics of a more general theoretical nature. Both categories of problem areas are considered next.

C.3.1 Critical Problem Areas

1) Higher-Dimensional Image Processing.

Past work in image acquisition, processing, and display has been conceived mainly with two-dimensional (2-D) images and image sequences in mind. However, rapid advances in computer and VLSI technologies in recent years have made it possible for researchers in MDSP to attack important and exciting problems involving higher-dimensional data. Examples of such problems include the analysis and visualization of dynamic CT and MRI volume data (which are scalar functions of the spatial and temporal variables, x, y, z, and t, respectively) for clinical applications, and of velocity fields in turbulent fluid flow (which are vector functions of x,y,z,t) to aid in the development of theories of turbulence. Challenging issues include: (1) How do we conquer the curse of dimensionality (not just extending lower-dimensional techniques in a brute-force manner)? (2) How do we display to a human viewer high-dimensional data in a meaningful way? Future work in MDSP needs to address 3-D and 4-D imaging, 3-D displays, and vector image processing (fluid flow, color, multispectral, multisensor).

2) Real-Time, Adaptive Image Acquisition and Processing

One of the major impediments to making many image processing algorithms operate in real time is the sheer amount of data that must be processed and the number of computational operations that must be performed. A potential solution to this problem lies in the development of algorithms and methods of data acquisition that are selectively adaptive. A natural example is the human eye which has high spatial resolution in the fovea and lower resolution on the periphery, and which is steered so that the fovea images the region of greatest interest. The design of such cameras is technically feasible but whether such systems can perform vision tasks more efficiently than more traditional systems will require research not only in image processing but also in computer vision and control.

3) Content-Addressable Image Databases

Many applications, such as forensic science and interactive video, involve large databases of images onto which “content-based” indexing need to be applied. One example of this problem is a fingerprint database from which all the images where three connected loops occur near the tip are required for matching against a partial print. Content-addressable image databases are extremely challenging not only since they must provide some type of image feature description language for indexing, but also because the method which is chosen to compress and decompress the images should allow for efficient indexing into the database using this feature keying.

4) Simultaneous Model Identification and Image Formation/Restoration

In many imaging applications, motion or focus errors cause severe degradation of image quality. For example, involuntary patient motion, such as breathing, is a source of degradation known as "ghosts" in MRI. Motion errors even on the order of a wavelength totally degrade images in SAR. Motion and defocus results in blurring in optical imaging. Often the source of degradation is not known and needs to be identified from the actual degraded data. MDSP theory has played a significant role in the understandings of these problems in the last decade. However, much future work is needed to completely solve these types of simultaneous identification/restoration problems. Performance bounds on the identifiability of the degradation models and on the quality of the image formation/restoration results using the identified models are needed.

5) Image Quality Evaluation and Theoretical Bounds

Methods for image compression and image restoration have resulted in a wide variety of techniques whose performance is very difficult to compare. Quantities related to signal-to-noise ratios are often quoted, but these are usually not subjectively meaningful. There is a need for objective quality measures that can be readily applied to a decoded or restored image to assess its subjective quality. There is a similar, important need to determine and assess theoretical bounds on the performance of coding and restoration algorithms.

6) Image and Scene Modeling

For all tasks of image processing (coding, enhancement, restoration, analysis), the question of modeling is paramount, not only in algorithm design, but also for image quality evaluation. For any given problem, appropriate modeling is almost always the key to a satisfactory solution. Important research problems include image acquisition modeling and image content modeling. Another aspect of modeling is the automatic construction of static and dynamic 3-D scene models from 2-D images. This problem is arising in many diverse fields including computer vision, photogrammetry and multimedia systems. To devise realistic, yet manageable models for many 3-D objects/scenes (e.g., hair, clothing) is extremely difficult.

7) Compression

A single-frame image commonly used in a wide variety of applications typically requires between 1MB and 4MB for storage in raster form. In many applications a time sequence or image stack is needed, and new imaging systems may acquire vector information rather than scalar for each pixel. New compression methods are needed for efficient storage and transfer of these large volumes of data. These methods should take full advantage of multidimensional structures and correlations present in the image data and should not be only a simple extension of successful 1-D algorithms. Compression methods tuned to restricted classes of images or specific models of image objects or image content have the potential for greater compression factors. In some applications, such as clinical imaging, lossless compression may be desired. A low bit-rate, lossless compression technique based on model or image class information would be a major breakthrough.

8) Image Fusion

In many applications, information must be extracted from multiple images acquired under different conditions or from different modalities. For example, information from a low-resolution PET image may need to be interpreted in the context of a high-resolution CT image, or information from a stack image of a city may need to be augmented with information from a map image. Alignment will typically require scaling, translating, and rotating in two or three dimensions. The parameters for these transformations must be derived from the relationship of the contents of the images. Advances in image understanding will be needed to make good estimates of the geometric relationship between the images.

9) Symbolic Algorithm and Data Manipulation

As multidimensional signal processing stretches the limits of our computational abilities, the need both for "customized" algorithms and for focused manipulation increases. For customized algorithms, algorithms can be specified in their most general form and then, using a high-level compiler which incorporates signal-processing equivalence rules, more effective methods for completing the same computation can be provided. For example, by customizing a general modulated filter bank to an application requiring dense temporal sampling, the computational requirements can be reduced from O(NlogN) to O(N). Further reductions in computational requirements may also be possible by combining the traditional numerical methods of MDSP with symbolic information, such as may be available from later processing stages in an image understanding system. In particular, symbolic information (or hypotheses) about the image content can control and focus the numerical processing of the image data. This approach should also allow for more robust image analysis by providing a "feedback path" from the later "recognition" stages to the early signal processing stages.

C.3.2 Critical Needs in Theory

1) Nonlinear Processing

Nonlinear approaches to the processing of multidimensional signals have shown considerable potential. Three disciplines that have developed different identities are neural networks, morphology, and order-statistics based nonlinear filters, including stack filters. At this time, applications of these approaches to the solution of problems in MDSP is hampered by a burgeoning of proposed structures and relatively incomplete analytical characterization of the behavior of algorithms based on them. There is a need for simple, yet flexible structures that can be tailored to the applications at hand in a well-defined way.

2) Statistical Modeling and Processing

Statistical models provide an important mechanism for dealing with the uncertainty and complexity inherent in multidimensional signals, while focusing on essential features. However, traditional statistical models have difficulty with the inherent nonstationarity and existence of extended regular features in image data. With random field models that can overcome these limitations, one is faced with an extremely difficult task of estimating the model parameters. More work is needed to develop models that have the capability to incorporate nonstationarity and extended features while permitting simple estimation of model parameters.

3) Multiscale and Time-Frequency Analysis

The modeling and analysis of nonstationary signals and images is currently being approached through the use of wavelets, the Wigner distribution and short-time Fourier transform, and fractals. Further work on these topics and on new, promising multiscale approaches is needed. This subject area is discussed in more detail elsewhere in this report.

4) Optimization Techniques

Many problems in MDSP, such as inverse problems, system identification, and signal design, can be formulated as optimization problems. Convergence to the global optimum is highly desirable in numerical optimization problems. Techniques such as simulated annealing and genetic algorithms have recently been used in MDSP applications. Further work is needed in the evaluation of these algorithms for MDSP applications, and in the development of other algorithms that have the global convergence property and faster convergence. A particularly useful iterative algorithm that has been introduced into the MDSP area within the last decade is projection onto convex sets (POCS). Unlike true optimization algorithms, POCS finds a feasible solution that satisfies certain constraints. Future work on POCS may include incorporation of a cost function to choose among many feasible solutions, or extension to nonconvex sets.

5) Mathematics of Inverse Problems

Inverse problems commonly arise in signal processing applications within the physical sciences and engineering. Although much work has been done on inverse problems in the past, there is still room for more work, especially in the solution of certain integral equations and PDE's. Regularized inversion of time/space-varying systems is a particularly important area that can be addressed by using SVD or Kalman filtering techniques. Also important is the inversion of systems with random impulse responses.

6) Algebraic Methods

Historically, the analysis tools in signal processing were often Fourier and convolution based. In recent years these operations and many others have been profitably viewed within the framework of linear algebra, where the useful concepts of subspace, projection, etc.,. arise in a natural way. There is still need for further development and application of linear algebra in MDSP, where the signal of interest may be a function of several variables, or may be a vector quantity. As a first-order extension of linear concepts, there is need for consideration of polynomial systems and their analysis, perhaps through application of algebraic geometry.

C.4 INITIATIVES

There are two interdisciplinary research areas, involving considerable MDSP, where NSF might profitably launch initiatives in conjunction with other funding agencies. It can be expected that many technological spin-offs would result from research initiatives in these areas. The potential for applying image processing and computer vision in these areas is briefly discussed below.

- MULTIMEDIA

- INTELLIGENT VEHICLE HIGHWAY SYSTEMS (IVHS)

C.4.1 Multimedia

Multimedia can be defined as communication between human and computer, and between humans via computer, which contains the following 3 elements:

(i) The communications between human and computer involves several media (e.g., text, still images, video, speech), and multiple senses (e.g., vision, hearing, tactile/force).

(ii) The communication is interactive (virtual reality, possibly).

(iii) The database is structured such that retrieval is very flexible.

Multimedia in its current fledgling state of development usually consists of coordinated interactively driven presentations of video sequences, audio, computer generated animations and still frames. As an example, a laser disk used for educational purposes typically will hold 120 to 200 minutes of analog “live video” plus the audio and indexing information to allow user-driven retrieval sequences. (Lucas Arts, Turner Broadcasting).

Potential applications of multimedia abound in training and education, business, the office environment, and in entertainment. Multimedia presentations have great potential for improving educational breadth and depth in all subjects as well as professional and commercial training and on-site problem solving assistance. One of the limiting factors to achievement of this potential is the ability to gather, store and efficiently index large amounts of video data. This will require efficient multidimensional compression for digital image sequences as well as advances in image understanding that will facilitate location and cross referencing by image control. In addition, coordinating different image databases of the same objects requires advances in image fusion. There are also technically challenging issues in how the human can interface with the computer, in terms of both computer input and display. For full realism, it will be necessary to develop 3-D displays with extraordinarily high resolution.

Multimedia is still in its infancy, so it is difficult to fully appreciate its potential and to project the forms that it might take. It is likely, though, that with the increasing power of computers and higher density of memory, and with the rapid progress in signal processing and computer vision, we are on the brink of a revolution. Unquestionably, this is an area of immense economic opportunity for the United States.

C.4.2 IVHS Initiative

There is a mounting research thrust in Europe, Japan, and the United States on the topic of Intelligent Vehicle Highway Systems (IVHS). The envisioned role of image processing and computer vision in these systems has been extremely ambitious (e.g., totally autonomous vehicles) and, therefore, unrealistic for the foreseeable future. It is apparent that huge sums of money will be spent on IVHS over the next 20 years, since it is expected to be more economical and environmentally responsible to efficiently use the roadway that we currently have, rather than build new highways. If properly used, image processing and computer vision holds great potential to become an important component of IVHS. Therefore, the National Science Foundation in conjunction with other government agencies may wish to consider a special initiative in image processing and computer vision for IVHS. Numerous research topics and approaches might be considered in this initiative. For example, aerial SAR could be used (even in inclement weather) to monitor traffic in metropolitan areas, with the SAR imagery automatically analyzed, and then appropriate data communicated to drivers to aid them in route selection. SAR could also be used on an automobile to present high-resolution all-weather imagery to the driver for improved safety. Various displays might be considered, including Heads-Up Displays (HUDs). Optical cameras and image processors could be installed at intersections to automatically monitor traffic flow and adjust traffic lights accordingly. Optical cameras, and sensors using other wavelengths, could be used for real-time monitoring of the state of the road surface. These are but a few illustrative topics that might be pursued under an IVHS initiative in image processing and computer vision.

APPENDIX D DIGITAL REPRESENTATION REPORT

This report summarizes the discussions, conclusions, and recommendations of the Digital Representation Group of the NSF Signal Processing Workshop held at Keystone, Colorado, 15-17 November, 1991. The group members were Robert M. Gray (Chair), Phil Chou, N. Jayant, Tom Lookabaugh, Steve Townes, Don Tufts, and Martin Vetterli.

The first morning began with brief self-introductions by all group members, each describing their own past and current research or application interests along with an occasional statement of general philosophy. A rough agenda was developed for the first day:

1) find a sentence or two describing what the group understood “Digital Representation” to mean;

2) describe in general terms and illustrate by example the role of digital representations in signal processing systems. The general examples can illustrate techniques and methods that have proved important in the past and can suggest important future avenues of research and important applications areas.

The agenda for the second day was:

3) describe specific research areas that are important for fundamental research and for applications in the long and short term;

4) conclude with specific recommendations of possible NSF initiatives and other actions.

Other specific topics were introduced and discussed, including the effect of standards on research and educational concerns. This report also contains a collection of miscellaneous comments made by the participants that did not clearly fit into the formal agenda.

D.1 Digital Representation

A digital representation can be the result of the conversion of an input signal into digital data or the starting point for the synthesis of an output signal or both. The general goal of such conversion is to provide a representation that is efficient and effective for a specific application.

Here the word “signal” is taken in the general sense of a possibly multidimensional waveform, time series, vector, or matrix. Efficient can mean, for example,

- low cost,
- low processing delay,
- low power,
- simple, or
- suitable for reduced subsequent signal processing.

Effective can mean, for example,

- high quality or that
- the digital representation preserves the most important features (or the necessary information) for the target application.

The meanings of both “efficient” and “effective” depend strongly on the specific application.

D.1.1 Digital Representations in Signal Processing

A general signal processing system involving digital representations either converts an input signal or some portion thereof into digital data or converts digital data into an output signal or both. The conversion can be direct or involve multiple steps, as will be considered shortly. The digital data itself can then be stored in a digital medium and later retrieved or communicated on a digital link. The retrieved or received data is delivered to a user for a specific application, several examples of which will be mentioned shortly. The key facet distinguishing the focus of “digital representation” is that at some point in the system the signal is described by digital data in a way suitable for the application. The group also considered issues relating to signal representation in general where such representation is important for digital representation. For example, purely continuous representations can play an important role in finding a good digital representation. What this representation should be and how it is formed usually depends on the specific class of signals and the application.

Common signal classes include:

audio
speech
images
  still frame images
  video
  character (e.g., bitmap)

graphical
multi-spectral (medical and aerial)
dynamic descriptions of handwritten documents
sensor data
geophysical data, e.g., seismic data
radar
sonar
infrared
medical (e.g., EKG, EEG)
multisensor (distributed)
multimedia
communications data

Common applications enabled by the digital representations of the signal include

data analysis
detection, parameter estimation, data fusion
classification
recognition
enhancement
diagnosis (e.g., radiology)
entertainment
digital audio, television, and multimedia
communication (transmission or storage and reconstruction or retrieval)
information services, databases
visualization (e.g., displaying 3-D information in useful ways)
synthesis (e.g., text to speech, text to images)
digital process control
automotive electronics
digital signal processing for location and control
metrology
security

In many applications, the digital representation is an intermediate form which separates conversion processes. In others, the digital representation is either the final result (in pure analysis) or the starting point (in pure synthesis).

Our focus is the system between the signal source and the digital representation. This is primarily because the synthesis portion is often either embedded in the analysis operation or is effectively an approximate inverse to the analysis. Some purely synthesis examples will, however, be considered later.

D.1.2 Operations Generating Digital Representations

We considered four basic types of common operations acting on signals and producing digital representations. These are described with a strong caveat: this is a rough clustering of techniques, and there are overlaps and natural combinations not easily separated. Most techniques can be used in fixed or adaptive ways. In addition, we are only describing a point-to-point system. The role of digital representation becomes more complicated in a distributed environment where representations may be spread among physical locations and the division of signal processing tasks are less clear. All of the operations described may be done in multistage or hierarchical ways. Those digital representations that are "good" may be quite different in such alternative architectures. For example, there may be an initial, conventional ADC of an input signal followed by a special representation which removes or separates certain components of the signal such as impulsive components. This representation may be followed by yet another representation such as the digital values of the model parameters of the preserved components.

Basic Conversion Operations:

1) analog-to-digital conversion
2) compression
3) decomposition
4) modeling

Most conversion systems can be broken up into the above pieces, which are described below. Clearly the division is not always obvious, and some operations may not naturally fit into this division. In all cases, both prescribed and adaptive versions can be used.

The ADC and compression operations are often grouped together under the general heading of "source coding."

1) Analog-to-Digital Conversion (ADC)

This is the conversion of an analog signal into a discrete-time sequence of digital data. In its basic form, this usually means a direct conversion of an analog waveform into a digital approximation.

Implementation issues dominate industrial ADC development. There is strong demand for high-quality, fast converters in the telecommunications and entertainment industries. Digital cellular phones, CDs, DATs, and the future enhanced and high-definition TVs promise high volume for high-quality, low-cost converters.

Some important signal processing issues are:

- how to best trade off the analog versus the digital processing in ADC;
- how to best make ADC robust against variations in source behavior and channels;
- how to match ADC to specific applications. A related issue is how to use ADC output bit streams directly in subsequent signal processing without an interim DAC. An example is performing transforms using the binary output of an oversampled ADC;
- how best to incorporate adaptive signal processing into ADC;
- how to use signal processing to improve ADC performance (e.g., using dither, compensation, and adaptation to vary input behavior);
- how to optimize pre- and post-filtering so as to complement the conversion.

2) Compression

Compression is the conversion of an analog or digital signal into a new digital signal having a reduced bit rate. This can be lossless (e.g., entropy coding) or lossy (e.g., transform coding).

The traditional heart of compression is the wide variety of compression algorithms that are currently in use or under study. Many of these have been standardized, a fact which led to the discussion of standards and research summarized later. The most common purpose of compression is for efficient storage and communication, especially of images, speech, and measurement data. No consensus was attempted for the relative merits of the numerous algorithms. Instead, important aspects were considered that have not yet received sufficient study.

- Significant improvement of existing compression algorithms for images, audio, and speech is likely to require more and better use of perceptual modeling in order to better reproduce perceptually important information at the expense of unimportant information. This field involves finding simple models that can be incorporated into compression algorithms, either by bit allocation strategies, weighted distortion measures, or transforming the original input into a “perceptual space,” e.g., filtering to enhance important spectral coefficients and downplay others.

- Compression must be made more robust against channel variations. This will be particularly important for packet network applications, where compression algorithms and strategies for reconstruction will have to take packet loss into account.

- Compression must be made more robust against source variation in cases where the compression system must match the source statistics.

- Pre- and post-processing should be better matched to compression so as to provide the best overall performance.

- Multiresolution techniques are increasingly important for progressive transmission, for adjustable rate applications, and for providing useful data structures for subsequent applications such as recognition and classification. These techniques are relatively young and include pyramid codes, wavelet codes, subband codes, and tree-structured codes.

- Special techniques should be developed for compressing distributed data that will subsequently be fused.

3) Decomposition

This is the decomposition of an original signal into simple or useful components. The decomposition can be achieved by transforming the original signal into analog parameters, as in Fourier analysis, principal components analysis, or wavelet representations. In this case, a digital representation is achieved by subsequent ADC or compression. The decomposition can be directly into discrete components, as in a classifier which labels signal segments.

Signal decompositions often can improve application performance and simplify subsequent signal processing such as quantization. It can include simple techniques such as well understood linear transforms or more sophisticated techniques such as the decomposition into multiresolution representations previously mentioned. Such decompositions are often immediately useful for providing digital representations preserving essential information, but they also can provide a very useful data structure for a variety of signal processing applications. On a related note, such decompositions can provide a useful theoretical framework for analyzing signal processing systems.

Two key attributes of decomposition methods are simplicity and the value of the decomposed signal for the subsequent application.

4) Modeling

Modeling is the construction of a mathematical model to fit the input signal. Models take two forms: models that produce the original signal (at least approximately) and models for the underlying process that generated the signal. The first class attempts to faithfully reproduce the complete signal and includes wire frame and facet models for images. Iterative maps or fractals can also provide such models for signals such as images as fixed points of the map. Common models of the second class include parametric models such as Gaussian and Markov models. Models that fall into both categories are those that approximate the original signal and then model the residual stochastically. Such models include the linear predictive models, neural nets, and classification and regression trees. As with decomposition, modeling can be done in a purely analog fashion and the subsequent parameters quantized, or the modeling can involve the direct selection of the “best” model for an input signal from a finite class, thereby producing a digital representation directly.

Popular examples of models for signal production are linear predictive models, articulatory models for speech, hidden Markov models, models of motion for motion compensation, neural nets, array manifolds, mathematical morphology, trees, and iterative maps.

Models can serve a variety of functions. They can be used to adapt signal processing to the signal, they can replace the signal itself as the object used by the application, and they can be used to select and extract important features from the signal.

Some issues arising in modeling follow.

- How do modeling and decomposition operations best interact?
- What features or attributes of a signal should be modeled to best perform a specific application? That is, what sort of preprocessing, linear or nonlinear, should be a front end to modeling?
- What are good classes of nonparametric models that allow one to be arbitrarily accurate in exchange for increasing complexity? For example, neural nets and decision trees can provide ever more accurate models at ever higher complexity. How are these models best designed?
- Can learning theory help us understand how much learning data is needed to construct trustworthy models?

D.2 Important Research Problems

The previous description of the four components of conversion into digital representation listed several issues and questions that the group thought important. In this section we identify several specific research areas, both theoretical and practical, that we believe merit increased effort and support. The order of the areas is not important.

Experimental:

Perceptual Models. This includes hard-core vision and audition studies in order to better model the human perceptual system. It also involves the soft-core development of techniques to design “perceptually optimal” representations based on knowledge of perceptual factors. This work is experimental and would best be done as an interdisciplinary study involving signal processing and psychology groups in both academia and industry.

Model Verification and Extension for Data Analysis. For example, is sea clutter best modeled as stochastic or chaotic?

Determining the Effects of Phase Properties of Filter Bank Decompositions on Perception.

Theory:

Determine supporting theory for models and decompositions found useful in practice.

Describe the behavior of signal decompositions such as quadrature mirror filter banks and wavelet decompositions in the presence of quantization noise and aliasing.

Mathematically quantify the tradeoffs between sampling rate and bits per sample in compression systems with an overall bit rate constraint. More generally, when extracting functionals of the data, how should one trade off the number of functionals with the bits for each in order to optimize performance for a specific application?

Performance analyses of adaptive signal decompositions. How well do adaptive techniques such as adaptive principal components analysis work for small sample sizes, and how does this affect the quality and usefulness of the subsequent digital representations?

Incorporate reasonable complexity measures into Shannon's rate-distortion theory. For example, complexity interpreted as power use can be commercially important. The goal is to quantify the achievable performance when constraints other than bit rate are considered.

Develop perceptually meaningful distortion measures so that Shannon's rate-distortion theory can be better applied to image and audio systems. The “perceptual entropy” of images can be roughly determined--that is, the entropy of a coded image at which the reproduction is perceptually indistinguishable from the original. Nothing is known, however, about the "perceptual rate-distortion function" providing the best bit rates when a certain perceptual error is permitted. The appropriate rate-distortion theory will likely have to consider other measures of performance than average distortion. This problem will require both theoretical and experimental work and cooperation among signal processors and psychologists.

Quantization noise analysis. Exact analyses exist for relatively few and idealized systems including quantizers, and most derivations use standard linearized approximations. Using nonlinear systems techniques to better characterize the noise behavior can provide important design insights and may eventually yield better systems. This is particularly true of oversampled ADCs and digital control systems with quantizers in feedback loops.

How can multistage systems be better designed by taking interactions between steps into account rather than separately designing each stage? A specific example is the joint design of automatic gain control and ADC. More generally, how should resources such as bits or flops be allocated in multistage systems (such as multistage or progressive compression or data analysis algorithms)? Such systems include a sequence of representations. Similarly, a design methodology should be found for hierarchical or layered systems that takes into account the interactions between layers.

A related issue is how to determine when a digital representation is "good enough" for further signal processing past the immediate application. For example, can a digital TV signal be used not only for viewing but also for editing? A digital representation that is good for the immediate use may be poor for later stages. In some cases it needs to serve as a "sufficient statistic" for later processing as well as for immediate display.

Joint source and channel coding, especially in packet networks. How to design conversion systems that are robust against channel errors and distortion including bit errors, burst errors, fading, and packet loss? What are the fundamental performance limits?

D.3 Random Topics

D.3.1 Standards

With the adoption of ever more standards, there is a concern that they can have a negative effect on research, discouraging work on novel approaches to problems with standardized solutions. It was generally felt that standards do not “kill” research, but they do and should discourage large effort being spent on minor improvements, unless there is some alternative motivation such as a better understanding of underlying fundamentals. Few (if any) want to see Lenna at .5 bpp at quality and complexity close to JPEG. Genuine and significant improvements in performance, cost, or complexity will eventually find an audience. As a graphical depiction, Tom Lookabaugh plotted a curve showing the evolution of a standard with cost and time as axes. The standard's cost decreases slowly with increasing time. Points close to the standards curve would be of little interest to industry, but points well removed from the curve in either direction could be important--e.g., much better performance or much lower cost. As an example, software video compression with existing hardware would be significantly cheaper than current proposed standards and hence would have a potential market in computer video.

D.3.2 University/Industrial Interactions

There was a general feeling for continuing to stress university/industrial cooperation on research as being beneficial to both. The suggestion was made that workshops be held regularly with interdisciplinary representation from interested university and industrial groups to provide a means of continuing exchange.

D.3.3 Tools

There was general support for a software consortium which would provide public domain commented software for new and old signal processing algorithms, including standards such as p*64, JPEG, JBIG, and MPEG, G.721, G.722, AV.259, as well as competing algorithms for compression, decomposition, modeling, etc. A database of training and test signals for common use would be an aide to research and comparisons of techniques.

D.3.4 Education

There was concern that ABET (Accreditation Board for Engineering and Technology) no longer reflects educational needs for research in signal processing. The group felt that undergraduate requirements should be modified to reflect the need for a course in fundamental DSP theory with a supporting lab and that undergraduates should encounter the fundamentals of quantization, likely either in a digital filtering course (coefficient roundoff) or in a basic communications system course (PCM).

The group recommended that the NSF provide equipment grants for courses and for research projects in multimedia signal processing. Courses and projects should involve both undergraduate and graduate engineering students.

The group favored encouraging graduate students in signal processing to take courses in other departments, especially Computer Science, Statistics, Psychology, and Applied Math. Research projects will be stronger if they are interdisciplinary and involve signal processing combined with ideas from these other areas. Many of the most important new algorithms draw heavily from such areas, and the earlier students encounter them, the better.

D.4 General Comments

It is important to not consider individual signal processing operations in isolation. For example, the effects of a communication channel can influence the design of a source coding system. This is true in the classical point-to-point systems, but it is particularly true in packet networks. Joint source and channel coding is an important area. ADC should be matched to specific sensors. Where possible, the data acquisition system should be designed with the final application in mind.

Since the specific signal and application have a strong effect on optimizing a system, it is important to develop methodologies for designing good systems involving digital representations.

Novelty is the key in university research, but the product is king in industry. It is important to learn to recognize research results and promote research efforts that can provide a real advantage to an application product. Perhaps a separate research funding apparatus should be developed for university research that focuses more on short- and long-term development of new techniques that can be transferred to industry for implementation.

The general digital representation area is considered in different forms by several communities, including communications, signal processing, control, and computer science.

What is “best” depends strongly on the implementation, especially if it is in hardware or software.

EE signal processors have often done badly in the PR game and have lost out to other fields in the press (and hence, as a result, sometimes with Congress).

American academic engineering has a reputation for publishing theories, while the Japanese turn those theories into products.

D.5 Recommendations

Much of this report has consisted of implicit or explicit recommendations of research areas meriting support. These will not all be repeated here. Instead, some of the rhetorical questions posed during the early stages of the workshop are answered, and some particularly important areas and applications are described.

D.5.1 Short versus Long Term

Most of the research issues described so far can be considered short-term in the sense that they are basically active, mature, and healthy areas in need of specific improvements and refinements. Most posed problems will have a short-term benefit in terms of improving basic knowledge and improving application products. An area of increasing importance that is far less well understood is that of man-machine communication. It was felt that special long-term emphasis should be placed on promising novel approaches to the related areas of speech and image recognition and understanding, to speech and image synthesis (conversion of text into voice or an image described by the text), and to the development of new complex models for handling such complicated problems. Such models may incorporate signal processing with symbolic representations developed by computer scientists. They will need to be useful for multiple applications.

D.5.2 “Hot” Applications

Certain applications were felt to be of unusual importance to U.S. industry as well as being likely to benefit from advances in signal processing. It was felt that proposals for research yielding short-term benefits to important applications, especially if they are low-cost and high-volume, should be encouraged and viewed favorably in comparison to purely theoretical research. Any proposal should be strengthened by well thought-out explanations of the potential benefits of the research to specific important applications.

The following list closely resembles that of the earlier part of the report but should be considered here as describing complete systems, not just users of digital representations. The list is certainly not exhaustive but provides those applications deemed most important to the group.

- entertainment: digital audio, TV, and multimedia
- medical image pressing (PACS)
- medical data analysis
- information services, databases
- personal communications
- metrology
- nondestructive testing
- manufacturing processes/digital process control
- automotive electronics
- sensor (and other) data transmission, storage, reconstruction
- voice and image synthesis
- system identification (for adaptation, adaptive equalization)
- digital representations that support editing and composition
- remote education (HDTV)
- environmental monitoring
- voice mail

D.5.3 Potential NSF Initiatives

Three application areas stood out as being particularly important for U.S. industry. All three are commercially important, can significantly benefit from advances in signal processing, and provide excellent training in a broad range of signal processing and related disciplines.

- MULTIMEDIA
    - Systems combining voice, video, images, and data involve virtually the entire range of signal processing. An immediate benefit would be the ability to "telecommute"--to perform many work functions from a terminal rather than an office. Although FAX and terminals or personal computers combined with modems are a step in this direction, the bandwidth and capabilities are minuscule compared to what is often available at the workplace. This area is not yet bound by standards, so there is great freedom in choosing approaches.
- PERSONAL COMMUNICATIONS
    - Communications to the person, not to the location. This area is presaged by mobile telephone and provides a wide range of

interesting technical problems, including source and channel coding, video and voice compression, spread spectrum, and networking.

- ENTERTAINMENT VIDEO AND AUDIO, ADVANCED TV

    - Technically challenging and commercially blossoming.

APPENDIX E IMPLEMENTATION REPORT

This report summarizes the discussions, conclusions, and recommendations of the Implementation Group of the NSF Signal Processing Workshop held at Keystone, Colorado, 15-17 November, 1991. The group members were J. Allen, T. P. Barnwell, R. W. Broderson, L. H. Jamieson, T. Meng, M. G. Slaney, and E. Swartzlander.

The implementation group began with a free-form discussion addressing issues that have an impact on current and future signal processing research. From the free form discussion, the group was able to develop a flowchart in Figure E.1 relating four key areas of algorithms, architectures, implementation technology, and applications.

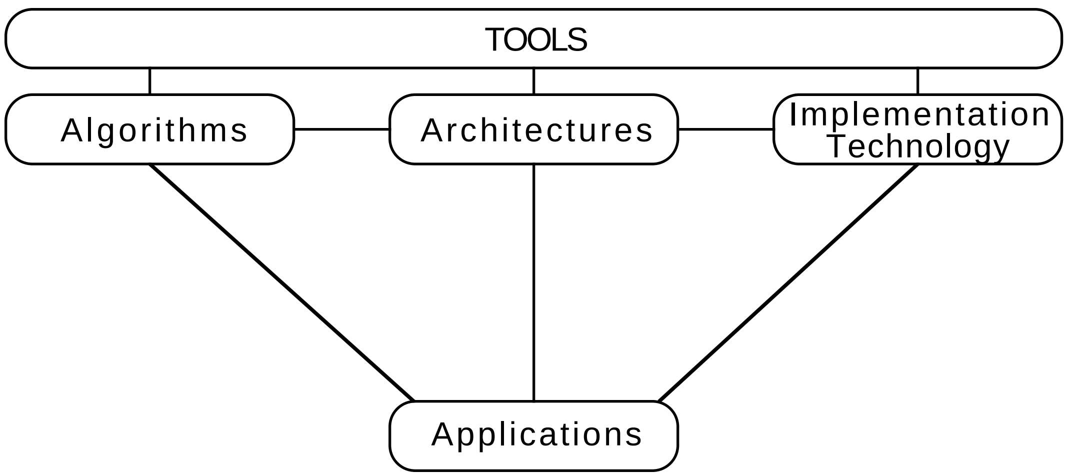

Figure E.1

During the free-form discussion, members of the group brought out the following items. It is necessary to define the application clearly before performing design. This provides a context for fine tuning the algorithm, which enhances the cost effectiveness of the resulting design. Taking an opposing viewpoint, certain colleagues felt that fine tuning algorithms is not always advantageous and that prototyping is necessary to uncover many “real world” problems that arise in any design.

Another item that needed to be considered is architecture exploration. In other words, how do you take an algorithm that is biased to some architecture and manipulate it to get a sense of the application space for that algorithm. Also what can we do to make it easier to try to perform the algorithm on alternative architectures? Committee members also felt that it was important to address the areas of integration constraints, parallelism between specific signal processing processors and general purpose processors, and areas which foster reusability of both software and hardware modules.

The concluding remarks that were made during the free-form discussion include the following items. It is important to define the extent to which the signal processing implementations are different from general purpose implementations and to define the class of applications that benefit most from specialized processors. The committee feels that it is necessary to define the audience and identify problems where implementations on specialized processors would be significant. This definition includes applications which require heavy duty usage and then show that the specialized signal processor is a necessary machine.

The special nature of DSP algorithms is both a problem and an opportunity. The problem is that they are often not well addressed by general parallel processing techniques. The opportunity is that their special nature makes it possible to use abstractions and techniques that cannot be effectively used elsewhere. Thus, in the signal processing area, special attention should be given to implementation efforts which seek to achieve high performance using the special properties of DSP algorithms.

E.1 Motivation: Major Applications

Signal Processing is a pervasive part of our world. In some cases, such as consumer CD players, the signal processing techniques used to convert the bit stream into an analog waveform are part of the marketing message used in advertisements. But more commonly, as in the example of DSP chips that control the position of a disk drive head, signal processing is a hidden but very important part of our lives.

Rapid advancement in our ability to implement DSP algorithms has made many new applications possible. Five years ago we understood the signal processing techniques that might be used to build a 9600 baud telephone modem, but the implementations were not yet common. With the advent of fast, cheap DSP chips, such modems are now widely available. Other applications of DSP that have only recently become possible due to new implementation techniques include, for example,

- ultrasonic doppler for cardiac blood flow imaging,
- cellular phones,
- synthetic aperture radar,
- automobile sensors (e.g., using exhaust gas temperature sensors to adjust engine parameters).

Clearly, faster, cheaper, and lower power technologies can drive signal processing algorithms and applications.

In the future, better implementation techniques for signal processing applications will have an even larger impact on our lives. For example, there are many new applications that will be simpler when robust speech recognition technology is available. Speech recognition research over the years has struggled to find algorithms that can be implemented in a reasonable amount of time given the current technology. While more algorithm work is important, it is faster implementation techniques that will make a large number of new algorithms available.

Another important use of signal processing in the future will be High Definition TV, which will require unprecedented signal processing power in a consumer-priced box to compress and decompress images so that they can be received, decoded, and stored. Other technologies, such as active noise control and better speech coding techniques, will improve the quality of our lives and make our world more efficient.

We have seen that there is a rich variety of important signal processing applications. They are pervasive, highly sophisticated, and complex and are central to national competitiveness. Rather than attack each of these many applications ad seriatim and by ad hoc methods, we advocate a two-pronged approach.

a) First, a framework for signal processing system design is to be developed. This structure provides an environment for a unified and interrelated set of signal processing tools ranging from high-level functional specification to detailed physical implementation of the technology.

b) In addition, we periodically exercise the design tools with an application that has significant signal processing challenges. A major task is to provide the overall system design for this facility. In addition, several signal processing subtasks are identified that constitute important research topics. In this way, a difficult but very attractive task provides an application context for state-of-the-art signal processing research.

Pursuit of the design of the application and the signal processing design framework will go on in parallel, and each project will motivate and inform the other. Practical experience from the design of the specific system will stimulate construction of the signal processing design framework, and conversely the design framework provides an overall perspective and emerging set of tools.

E.2 Framework for Performance-Directed Implementation Exploration

In order to solve the myriad signal processing research and development problems, it is essential that we develop a framework that supports the exploration of alternative implementations. Figure E.2 shows the structure of such a framework. The goal is to provide a signal processing user or signal processing system designer with tools to support a full spectrum of interactive or automatic exploration capabilities, including:

- user-suggested alternative algorithms
- automatically generated (e.g., by transformation or library look-up) alternative algorithms
- user-directed alternative architectures
- automatically generated alternative architectures
- software or hardware targets: software simulation on available serial or parallel architectures, hardware implementation in a target technology.

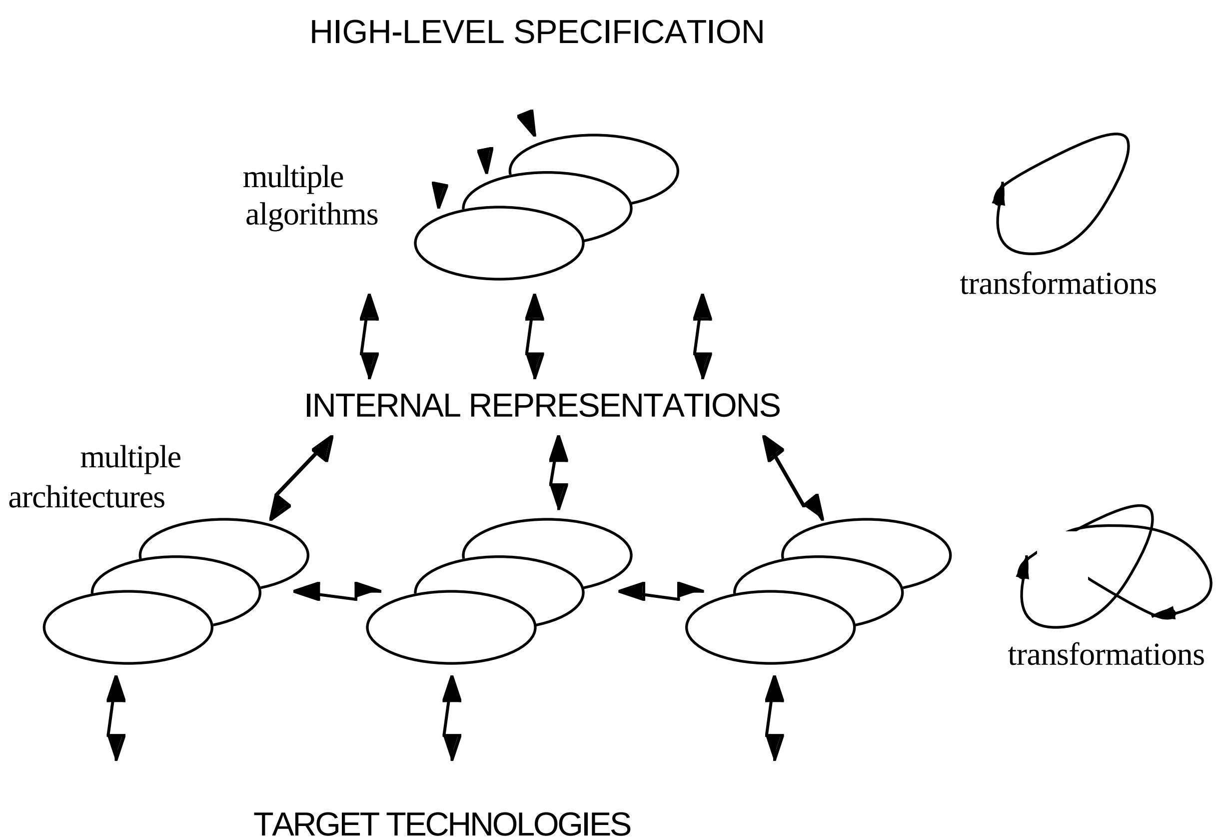

Figure E.2

At the highest level, the user specifies functional requirements and performance requirements. The bottom level represents the target hardware or software technology, or candidate technologies, from which implementation constraints are derived. The process of moving from the high-level specification to the target technology consists of exploring the algorithm and architecture spaces and includes transformations in the dimensions of parallelism degree and structure, interconnection topology, granularity, allocation, scheduling, data precision and representation, and component complexity.

Two aspects are key to the usefulness and feasibility of this framework:

1. The alternatives can be evaluated in a wide range of user-specified performance spaces, including processing rate, area, power consumption, cost, etc. Performance evaluation drives the exploration of the full algorithm architecture-technology space. For example, evaluation of the performance at the architecture level based on a target technology can result in acceptance of the implementation, can trigger the consideration of an alternative architecture, or can trigger the consideration of a different algorithm.

2. Three properties of signal processing algorithms are essential to this view of implementation exploration: (l) real-time execution requirements require the use of the performance-driven design paradigm that forms the basis of this framework; (2) the extraordinarily high degree of parallelism in signal processing algorithms and data provides the basis for the large design space that must be explored; (3) the prevalence of regular and canonical structures in signal processing algorithms permits the definition of transformations and libraries that make this otherwise complex design framework feasible.

E.3 Research Problems

Digital signal processing algorithms are a special subset of the class of algorithms that might be addressed in an implementation research program. As previously noted, they are computationally intense and highly structured and often possess much more parallelism than do other algorithms. In addition, they often have much simpler control structures, and a large percentage of them can be represented using signal flow and other graph structures.

The special nature of signal processing algorithms represents both a problem and an opportunity. The problem is that such algorithms are often not well addressed by general implementation techniques and tools because their intrinsically broader view prevents them from taking advantage of the special characteristics of signal processing algorithms. The opportunity is that their special nature makes it possible to use abstractions and techniques that cannot be effectively used elsewhere. Thus, in the signal processing area, NSF should seek to fund efforts which are aimed at implementation issues which are largely unique to signal processing.

The evolution of implementation techniques will be of extreme importance in signal processing over the next ten years, and it is very important that funded projects in the implementation area contribute in a substantive way to this evolution. Direct research on fundamental implementation issues should be encouraged, as well as projects that involve a specific implementation and also have a broad impact (in terms of new algorithms, new CAD tools, new understandings of fundamental problems, etc.). Because research on implementations is often more expensive than pure algorithmic research, combined funding with programs with common implementation interests is highly recommended.

In a broad sense, research should be supported which leads to or in some way supports the development of signal processing system design tools. These design tools should specifically seek to overcome the barriers which traditionally separate algorithm, architecture, and technological pursuits. Because of the immediate importance of signal processing to many different application areas and the universal need for rapid prototyping, there should be a specific emphasis on performance directed synthesis. In addition, the research should be performed in the context of a specific application.

A variety of research opportunities exist in combining algorithms, architectures, and technologies. Many of the critical research issues are concerned with representations, transformations, and mappings which allow design trade-offs to be effectively achieved. In general, these will be CAD tools and signal processing system compilation techniques that evolve in a common framework. Important topics include:

- new concurrent topologies and algorithms
- new, general algorithmic representational formalisms
- mapping techniques between different algorithmic structures
- automatic scheduling and resource allocation techniques (compilers) for parallel and/or deeply pipelined systems
- effective library generation and utilization techniques for algorithms, schedules, cells, software, etc.
- signal processing-oriented debugging and simulation tools.

E.4 Infrastructure Initiative

The implementation of a complete design system as outlined above will require a considerable effort from the community at large. Such a system consists of a common design infrastructure on which individual research tools can be developed and distributed. These tools will perform new transformations and explore scheduling and mapping strategies which convert signal processing algorithms into implementations. It will also support design principles such as real-time constraints and regular transformations which are characteristic of signal processing. The common infrastructure is important since it reduces the amount of duplication of effort, maximizes the sharing of tools, and provides a mechanism to evaluate the performance of these tools.

The common design infrastructure (Figure E.3) flows from a variety of entry forms, both graphical and textual, to implementations in any of several technologies including custom chips, boards, multi-component modules, programmable logic, and programmable DSP modules.

Three levels of abstraction are shown: dataflow, structural, and physical. The transformations between these levels and within these levels will be accomplished by tools developed in the research community. At each level there are simulators and means for viewing the design.

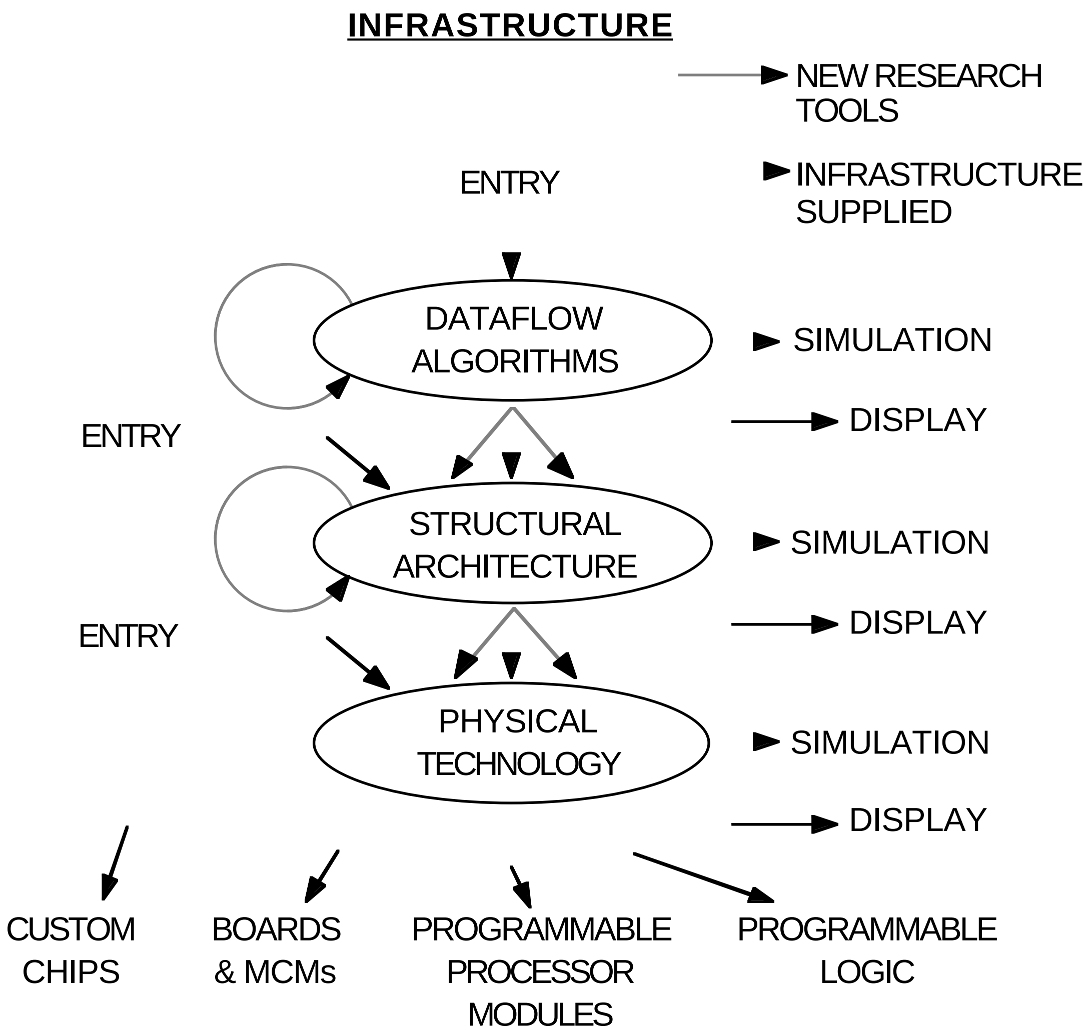

Figure E.3. The basic levels of abstraction and support for the infrastructure.

- SIGNAL PROCESSING INFRASTRUCTURE

  - The infrastructure is supported by a consistent underlying database format. The transformations, scheduling, etc., provided by the new research tools read the appropriate information from the database and then write their optimized results back into it. Though the definition of such a database standard has been the goal of the general-purpose CAD community for some time with only minimal success, it is believed to be possible in the more constrained domain of signal processing applications. This is particularly true since experience with such a database has been obtained at a number of locations who have been involved in signal processing CAD tools.

APPENDIX F LIST OF ACRONYMS

ABET Accreditation Board for Engineering and Technology
ADC analog-to-digital conversion
ASIC application-specific integrated circuit
ASP analog signal processing
ASSP Acoustics, Speech, and Signal Processing Society of IEEE

CAD computer-aided design
CAT computerized axial tomography
CCITT Comite Consultatif Internationale Telegraphique et Telephonique
CT computer tomography

DSP digital signal processing

EEG electroencephalogram
EKG electrocardiogram

GPS global positioning system

HDTV high-definition television

IEEE Institute of Electrical and Electronics Engineers
IQ iterative quadrature
IVHS intelligent vehicle highway systems

MCM multiple-chip module
MDSP multidimensional digital signal processing
MIMD multiple-input, multiple-data
MIPS microelectronics and information processing systems
MRI magnetic resonance imaging

NDT	nondestructive testing
NL	nonlinear
NMR	nuclear magnetic resonance
NSF	National Science Foundation
OSP	optical signal processing
PCM	pulse code modulation
PCN	personal communication network
PET	positron emission tomography
QMF	quadrature mirror filters
R&D	research and development
RF	radio frequency
SAR	synthetic aperture radar
SEM	scanning electron microscope
SIMD	single-input, multiple-data
SPECT	single photon emission computer tomography
VLSI	very large-scale integration
VSAT	very small aperture terminal